

06-27-00

PTO/SB/05 (12/97)

Please type a plus sign (+) inside this box → ☐

Approved for use through 09/30/00. OMB 0651-0032

Patent and Trademark Office: U.S. DEPARTMENT OF COMMERCE

Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it displays a valid OMB control number

06/26/00
jc862 U.S. PTO

06/26/00
jc862 U.S. PTO

UTILITY PATENT APPLICATION TRANSMITTAL <small>(Only for new nonprovisional applications under 37 CFR 1.53(b))</small>	Attorney Docket No.		TI-29338
	First Named Inventor or Application Identifier		Donald E. Steiss
	Title	Sub-Pipelined And Pipelined Instruction Execution In A Vliw	
	Express Mail Label No.		EL 547746403US

APPLICATION ELEMENTS <small>See MPEP Chapter 600 concerning utility patent application contents</small>		ADDRESS TO: Assistant Commissioner for Patents Box Patent Application Washington, DC 20231	
1. <input checked="" type="checkbox"/> *Fee Transmittal Form (e.g., PTO/SB/17) <small>(Submit an original, and a duplicate for fee processing)</small>		6. <input type="checkbox"/> Microfiche Computer Program (Appendix)	
2. <input checked="" type="checkbox"/> Specification <small>(preferred arrangement set forth below)</small> - Descriptive title of the Invention - Cross References to Related Applications - Statement Regarding Fed sponsored R&D - Reference to Microfiche Appendix - Background of the Invention - Brief Summary of the Invention - Brief Description of the Drawings (if filed) - Detailed Description - Claim(s) - Abstract of the Disclosure	[Total Pages] 15	7. <input type="checkbox"/> Nucleotide and/or Amino Acid Sequence Submission <small>(if applicable, all necessary)</small> a. <input type="checkbox"/> Computer Readable Copy b. <input type="checkbox"/> Paper Copy (identical to computer copy) c. <input type="checkbox"/> Statement verifying identical of above copies	
3. <input checked="" type="checkbox"/> Drawing(s) (35 USC d113)	[Total Sheets] 2	ACCOMPANYING APPLICATION PARTS 8. <input type="checkbox"/> Assignment Papers (cover sheet & Documents(s)) 9. <input type="checkbox"/> 37 CFR §3.73(b) Statement <input type="checkbox"/> Power of Attorney <small>(when there is an assignee)</small> 10. <input type="checkbox"/> English Translation Document (if applicable) 11. <input type="checkbox"/> Information Disclosure Statement (IDS)/PTO-1449 <input type="checkbox"/> Copies of IDS Citations 12. <input type="checkbox"/> Preliminary Amendment 13. <input checked="" type="checkbox"/> Return Receipt Postcard (MPEP 503) <small>(Should be specifically itemized)</small> 14. <input type="checkbox"/> *Small Entity Statement(s) <input type="checkbox"/> Statement filed in prior application <small>(PTO/SB/09-12)</small> Status still proper and desired 15. <input type="checkbox"/> Certified Copy of Priority Document(s) <small>if foreign priority is claimed</small> 16. <input type="checkbox"/> Other: <small>*A new statement is required to be entitled to pay small entity fees, except where one has been filed in a prior application and is being relied upon.</small>	
4. Oath or Declaration	[Total Pages]		
a. <input type="checkbox"/> Newly Executed (original or copy)			
b. <input type="checkbox"/> Copy from a prior application (37 CFR §1.63(d)) <small>(for continuation/divisional with Box 17 completed)</small>			
[Note Box 5 below] i. <input type="checkbox"/> DELETION OF INVENTOR(S) <small>Signed statement attached deleting inventor(s) named in the prior application, see 37 CFR §1.63(d)(2) and 1.33(b).</small>			
5. <input type="checkbox"/> Incorporation By Reference (useable if Box 4b is checked) <small>The entire disclosure of the prior application, from which a copy of the oath or declaration is supplied under Box 4b, is considered as being part of the disclosure of the accompanying application and is hereby incorporated by reference therein.</small>			

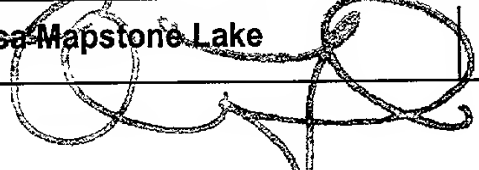
17. If a CONTINUING APPLICATION, check appropriate box and supply the requisite information below and in a preliminary amendment:

<input type="checkbox"/> Continuation	<input type="checkbox"/> Divisional	<input type="checkbox"/> Continuation-in-part (CIP)	of prior application No: /
Prior application information: Examiner		Group / Art Unit:	

18. CORRESPONDENCE ADDRESS

☐ Customer Number or Bar Code Label (Insert Customer No. or Attach bar code label here) or ☒ Correspondence address below

NAME	Rebecca Mapstone Lake Texas Instruments Incorporated		
ADDRESS	P. O. Box 655474, M/S 3999 Patent Activity		
CITY	Dallas	STATE	TX
COUNTRY	U.S.A.	TELEPHONE	(972)917-5645
		ZIP CODE	75265
		FAX	(972) 917-4408

Name (Print/Type)	Rebecca Mapstone Lake	Registration No. (Attorney/Agent)	36,208
Signature		Date	6.26.00

Burden Hour Statement: This form is estimated to take 0.2 hours to complete. Time will vary depending upon the needs of the individual case. Any comments on the amount of time you are required to complete this form should be sent to the Chief Information Officer, Patent and Trademark Office, Washington, DC 20231. DO NOT SEND FEES OR COMPLETED FORMS TO THIS ADDRESS. SEND TO: Assistant Commissioner for Patents, Box Patent Application, Washington, DC 20231.

SUB-PIPELINED AND PIPELINED EXECUTION IN A VLIW

Field of the Invention

This invention relates in general to instruction set architectures and more specifically to sub-pipelined and pipelined instruction execution in a VLIW.

Background of the Invention

Traditionally, Very Long Instruction Word, VLIW processors have been defined by the following set of attributes. The ability to specify multiple, independent operations in each instruction. (a MultiOp instruction.) VLIW architectures are horizontal machines, with each wide instruction-word or *MultiOp*, consisting of several operations or , *Ops*. All Ops in a MultiOp are issued in the same execution schedule. Programs that assume specific non-unit latencies for the operations and which, in fact, are only correct when those assumptions are true. The requirement for static, compile-time operation scheduling taking into account operation latencies and resource availability. Consequently, the requirement that the hardware conform exactly to the assumptions built into the program with regards to the number of functional units and the operation latencies. The absence of any interlock hardware, despite the fact that multiple, pipelined operations are being issued every cycle.

The original attraction of this style of architecture is its ability to exploit large amounts of instruction-level parallelism (ILP) with relatively simple and inexpensive control hardware. Whereas a number of VLIW products have been built which are capable of issuing six or more operations per cycle [4, 5, 3], it has just not proven feasible to build superscalar products with this level of ILP [18, 2, 14, 8, 7, 6]. Furthermore, the complete exposure to the compiler of the available hardware resources and the exact operation latencies permits highly optimized) schedules.

These very same properties have also led to the perception that VLIW processors are of limited interest as products. The rigid assumptions built into the program about the hardware are viewed as precluding object code compatibility

5 between processors built at different times with different technologies and, therefore,
having different latencies. Even in the context of a single processor, the need for the
compiler to schedule to a latency, that is fixed at compile-time, is problematic with
operations such as loads which can have high variability in their latency depending on
whether a cache hit or miss occurs. Because of this latter problem, VLIW products
10 have rarely adhered to the ideal of no interlock hardware, whatsoever. Alternatively,
with an implementation of a VLIW architecture, interlocking and stalling of the processor
is common when a load takes longer than expected.

The conventional wisdom is that dynamic scheduling is inapplicable to VLIW
processors. A first step towards understanding how to perform dynamic scheduling on
15 VLIW processors is to recognize the distinction between traditional VLIW processors
and the concept of a VLIW architecture. A VLIW processor is defined by a specific set
of resources (functional units, buses, etc.) and specific execution latencies with which
the various operations are executed. If a program for a VLIW processor is compiled
and scheduled assuming precisely those resources and latencies, it can be executed
20 on that processor in an instruction-level parallel fashion without any special control
logic. Conversely, a VLIW processor that has no special control logic can only correctly
execute those programs that are compiled with the correct resource and latency
assumptions. VLIW processors have traditionally been built with no special control
logic and this has led to the conclusion that VLIW processors must necessarily be
25 designed in this fashion.

A different view of VLIW is as an architecture, i.e., a contractual interface between
the class of programs that are written for the architecture and the set of processor
Implementations of that architecture. The usual view is that this contract is concerned
with the instruction format and the interpretation of the bits that constitute an Instruction.
30 But the contract goes further and it is these aspects of the contract that are of primary
importance in this patent. First, via its MultiOp capability, a VLIW architecture specifies
a set of operations that are guaranteed to be mutually independent (and which,

5 therefore, may be issued simultaneously without any checks being made by the issue hardware).

Second, via assertions about the operation latencies, an architecture specifies how a program is to be interpreted if one is to correctly understand the dependencies between operations. In the case of a sequential architecture, most latencies are
10 assumed by the programmer to be a single cycle. So, the input operands for an operation must appear to that programmer to be determined by all the operations that were issued (and, in most cases, completed) before the operation in question. This is true because there are some sequential architectures, SPARC for example, where some instructions (branches with delay slots) have non-unit latencies.

15 In the case of programs for VLIW architectures, with operations having non-unit latencies, the input operands for an operation are not determined by all the operations that were issued before the operation in question. What matters is the operations that are supposed to have *completed* before the issuance of the operation in question. Operations that were issued earlier, but which are not supposed to have completed as yet, do not impose a flow dependence upon the operation in question.

20 A program has unit assumed latencies (UAL) if the semantics of the program are correctly understood by assuming that all operations in one instruction complete before the next instruction is issued. A program has non-unit assumed latencies (NUAL) if at least one operation has a non-unit assumed latency, L , which is greater than one, i.e.,
25 the semantics of the program are correctly understood if exactly the next $L-1$ instructions are understood to have been issued before this operation completes. An architecture is UAL (NUAL) if the class of programs that it is supposed to execute are UAL (NUAL). We shall use the terms NUAL program and latency-cognizant program interchangeably.

30 Very Long Instruction Word, VLIW, processors are viewed as an attractive way of achieving instruction-level parallelism because of their ability to issue multiple operations per cycle with relatively simple control logic. Although VLIW architectures offer the advantages of simplicity of design and high issue rates, a major impediment to

5 the use of VLIW and other novel ILP architectures is that they are not compatible with
the existing software base. Lack of object code compatibility in VLIW architectures
across processors having different hardware latencies and varying levels of parallelism
is a severe limit to their adoption as a general purpose computing paradigm. This
means that an installed software base of binaries cannot be built around a family of
10 generations. The economic implications of this problem are enormous, and an efficient
solution is necessary if VLIW architectures are to succeed. Two classes of approaches
to this problem have been reported in the literature: hardware approaches and software
approaches. Although these techniques may provide compatibility, they do so at the
expense of hardware complexity that can potentially impact cycle time. A typical
15 software approach is to statically recompile the VLIW program from the object file. The
approach generates multiple executables, which poses difficulties for commercial copy
protection and system administration. For example, if a first generation of a machine
has certain latencies involved with each functional unit and the second generation
VLIW machine has different latencies involved with those same functional units, the old
20 VLIW schedule cannot be guaranteed to execute properly on the second generation
machine due to the flow dependence between the different operations. The same type
of problem results if the second generation machine includes an additional functional
unit. Even if the latencies remained the same, the code scheduled for this new
machine would not execute correctly on the older machines because the scheduler has
25 moved operations in order to take advantage of the additional multiplier. There is no
trivial way to adapt this schedule to the older machines. This is the case of downward
incompatibility between generations. In this situation, if different generations of the
machines share binaries, compatibility requires either a mechanism to adjust the
schedule or a different set of binaries for each generation. IBM describes hardware
30 features for an ILP machine called DAISY (Dynamically Architected Instruction Set from
Yorktown). DAISY is specifically intended to emulate existing architectures, so that all
existing software for an old architecture (including operating system kernel code) runs
without changes on the VLIW architecture. Each time a new fragment of code is

5 executed for the first time, the code is translated to VLIW primitives, parallelized and saved in a portion of main memory not visible to the old architecture, by a Virtual Machine Monitor (software) residing in read only memory. Subsequent executions of the same fragment do not require a translation (unless cast out). A limitation of the hardware approaches is that the scope for scheduling is limited to the window of Ops
10 seen at run-time, hence available ILP is relatively less than what can be exploited by a compiler. These schemes may also result in cycle time stretch, a phenomenon due to which many are considering the VLIW paradigm over superscalar for future generation machines.

15 Instruction set architecture is a description of a computer architecture at the level the programmer can observe. A *programmer's model* of a computer architecture is a similar term. In an exposed-pipeline architecture, the delays associated with pipelined execution of instructions are visible in the instruction set architecture and can be exploited to improve computational bandwidth.

20 Other techniques of solving the compatibility issues involve migrating the software across to a new machine architecture. In these applications, oftentimes the original old architecture that is trying to be emulated, or the native instruction set architecture of a processor, is called the base architecture. The VLIW which emulates the old architecture, or any instruction set architecture, other than the base architecture of a given processor that will be supported on that processor, is called the migrant
25 architecture.

Code compatibility between current and future generations of exposed pipeline VLIW DSPs is the example of compatibility reflected by this invention. For example, the TI C6000 DSP and the TI 64 bit C6000 DSP extension are the current and future architectures. The TI 64-bit C6000 DSP architecture contains changes in ISA visible
30 pipelining and other architectural features to provide higher operating frequencies. These changes are reconciled with requirements for binary compatibility and a robust code migration path implemented with the invention below.

Summary of the Invention

A subpipelined translation embodiment providing binary compatibility between current and future generations of DSPs is disclosed. When a fetch packet is retrieved from memory, the entire fetch packet is assigned an operating mode (base instruction set or migrant instruction set) according to the execution mode at the time the request was made to the instruction memory for the fetch packet. The fetch packets from the instruction memory are parsed into execute packets and sorted by execution unit (dispatched) in a datapath shared by both execution modes(base and migrant). Because the fetch packet syntax and execution unit encoding is different between the migrant and base architecture in this case, the two execution modes have separate control logic. Instructions from the dispatch datapath are decoded by either base architecture decode logic or the migrant architecture decode logic, depending on the execution mode bound to the parent fetch packet of the instructions being decoded. Code processed by the migrant and base decode pipelines produces machine words that control the register files and the execution hardware functional units. These machine words are selected with a multiplexer. The choice of the eventual machine word from the multiplexer is governed by the operating mode bound to the fetch packet that produced the machine word and sequencing logic for sub-pipelined execution. The selected machine word controls a global register file, which supplies operands to all hardware execution units and accepts results of all hardware execution units.

Brief Description of the Drawings

Figure 1 illustrates an example of migrant code of a subpipelined execution migrant architecture which has 1/4 the instruction latency the of the base architecture.

Figure 2 illustrates the rescheduling of the migrant code, running in subpipelined execution mode, after translation to the base architecture code, illustrating the inefficiency of the translation.

Figure 3 illustrates a sub-pipelined translation embodiment according to a preferred embodiment of the invention.

Detailed Description of the Preferred Embodiments

Subpipelined execution is a hardware-efficient method for executing code from a migrant exposed-pipeline architecture. In this method, the base architecture is designed with a fixed multiple(call this S) of the instruction latencies of the desired migrant architecture. With this relationship between the migrant and base architectures, code from the migrant architecture can be executed on the base architecture by delaying the issue of migrant instructions by $S-1$ clock cycles. In addition to providing the subpipelined execution mode, facilities are provided to change between base and migrant instruction sets with low overhead.

Given the code illustrated in Figure 1, and a base architecture with 4 times the instruction latency of the migrant architecture for all instructions, as illustrated in the table below:

Instruction	Migrant Delay Slots	Migrant Latency	Base Delay Slots	Base Latency
ADD	0	1	3	4
AND	0	1	3	4
OR	0	1	3	4
MUL	1	2	7	8

Table 1. Comparison between migrant and base architecture latencies

The code can be rescheduled in the base architecture as illustrated in Figure 2. As is illustrated, because of all the NOPs necessary to make up difference in latencies between the base and migrant codes, the code running in subpipelined execution mode will not exceed a fraction of the performance of the base architecture instruction set. Depending on the performance requirements and other characteristics of the code, two automated approaches are available to enhance the performance of code from the migrant architecture. First code from the migrant architecture can be re-linked with

5 libraries written in the base architecture, and second, off-line binary translation techniques can convert migrant architecture code to base architecture code.

In the case of migrant code with base libraries, the performance gain over the subpipelined case is proportional to the execution time spent in the base libraries and the performance of the library routines in the base ISA. The time spent in libraries
10 varies strongly with applications, but it is not uncommon for greater than 50% of the execution time of an application to be in libraries in the 32-bit C6000 DSP instruction set.

Off-line binary translation takes in binary code from the migrant architecture, reverse-assembles the code and converts it to the base architecture. The conversion
15 process can be implemented in several ways.. Conversion is simplified by having the sub-pipelined execution mode: cases that can be recognized as difficult or non-translatable execute in sub-pipelined mode. This allows off-line translation to be deployed incrementally: as the notion of difficult varies over time, the fraction of the code executing in subpipelined mode varies.

Although the previously discussed(in the Background of Invention) IBM and HP
20 and North Carolina State University solutions on supporting compatibility between generations of VLIW processors may result in VLIW base architectures that can support migrant architectures, these solutions are not appropriate solutions for DSPs. In particular, dynamic translation approaches as described by the IBM and NCSU works
25 fail to provide the run-time predictability required for real-time DSP applications. The method described in the HP work will have complex but deterministic run-time, but substantial hardware costs from the delayed -issue instruction buffer, delay register file and copyback unit.

Work has also been published involving CISC to superscalar RISC translation.
30 Later versions of this same work involves static translation coupled with software simulation of the migrant architecture running on the base architecture. In the present invention, all of the functions of the software simulation of the migrant architecture are replaced with hardware execution in subpipelined mode.

5 Figure 3 illustrates a subpipelined translation embodiment according to a preferred embodiment of the invention. In this invention, code is fetched from memory. The instruction memory may be implemented as a directly-addressed RAM or cache. The code in the instruction memory can appear to the programmer to be in the migrant instruction set or in the base instruction set. In an implementation, the code in the instruction memory may be pre-decoded to facilitate later instruction processing. A group of instructions simultaneously retrieved from memory is called a fetch packet. Fetch packets are assigned an operating mode (base instruction set or migrant instruction set) according to the execution mode at the time the request was made to the instruction memory for the fetch packet. This allows for rapid mode transitions.

15 The fetch packets from the instruction memory are parsed into execute packets and sorted by execution unit (dispatched) in a datapath shared by both execution modes(base and migrant). Because the fetch packet syntax and execution unit encoding is different between the migrant and base architecture in this case, the two execution modes have separate control logic.

20 Instructions from the dispatch datapath are decoded by either base architecture decode logic or the migrant architecture decode logic, depending on the execution mode bound to the parent fetch packet of the instructions being decoded. In the case of exposed-pipeline VLIW instruction sets, the decode logic for the base and migrant architectures primarily translates opcodes to the control signals required to execute the specified instructions on the execution hardware functional units. Because of the relationship defined in this invention with regard to the latencies of the migrant and base instruction sets(the base operations have twice the latency of the migrant operations) and the exposed pipeline characteristics of these instruction sets, instruction decoding techniques that need knowledge of the pipeline depth, instruction graduation and instruction dependencies are not required. This results in reduced hardware and complexity in the instruction decode logic.

30 Code processed by the migrant and base decode pipelines produces machine words that control the register files and the execution hardware functional units. These

5 machine words are selected with a multiplexer, which can also select a third No-operation instruction choice. The choice of the eventual machine word from the three possibilities is governed by the operating mode bound to the fetch packet that produced the machine word and sequencing logic for sub-pipelined execution.

10 In the first preferred embodiment of this invention, the selected machine word controls a global register file, which supplies operands to all hardware execution units and accepts results of all hardware execution units. Two hardware pipeline clock cycles later, the selected machine word controls local register files that supply operands to either local execution hardware functional units or neighbor hardware execution functional units. Finally, the selected machine word controls the various types of
15 execution hardware that evaluate functions on the operands to produce the results.

Claims

I claim:

1. The structure of a subpipelined translation embodiment providing binary compatibility between a base architecture and migrant architecture of a VLIW architecture comprising:

a VLIW architecture comprising a base architecture and a migrant architecture and having a base execution mode and a migrant execution mode;

a fetch packet retrieved from memory, the fetch packet having an operating mode in dependence upon the execution mode at the time the request was made to the memory for the fetch packet;

a shared datapath by both the base and migrant architectures for parsing said base architecture mode and migrant architecture mode fetch packets into execute packets and for dispatching those base execute packets to the appropriate base architecture decode of the execute hardware;

a migrant architecture control circuit for dispatching execute packet instructions having a migrant execution mode to a migrant architecture decode;

execute hardware for executing execute packet instructions on execution units and having a base architecture decode and a migrant architecture decode for decoding said base architecture instructions and said migrant architecture instructions, respectively, in dependence upon the execution mode of the fetch packet of the instructions being decoded, prior to executing;

a multiplexer having at least two inputs and one machine word output wherein one input is the output of said migrant architecture decode and the other input is the output of said base architecture decode, said multiplexer choosing in dependence upon the operating mode of said fetch packet;

machine words for controlling the execution hardware units.

2. The structure according to Claim 1, and further comprising a third input to said multiplexer wherein said third input is a no operation instruction

3. The structure according to Claim 1, wherein said machine word also controls registers.

10 4. The structure according to Claim 1, wherein said machine word controls a global register file, which supplies operands to all hardware execution units and accepts results of all hardware execution units.

15 5. The structure according to Claim 4, wherein said machine word controls local register files that supply operands to either local execution hardware functional units or neighbor hardware execution functional units subsequent to said machine word controlling said global register file.

20 6. The structure according Claim 5, wherein said machine word controls the various types of execution hardware that evaluate functions on the operands to produce the results of said hardware execution units subsequent to said machine word controlling said local register files.

25 7. The structure according to Claim 1, wherein the base and migrant architecture decode units translates opcodes to the control signals required to execute the specified instructions on the execution hardware functional units.

8. The structure according to claim 1, and further comprising said migrant architecture control circuit for issuing no-operation instruction to preserve the semantics of the instruction in the migrant architecture.

30 9. The structure according to Claim 1, wherein said VLIW architecture is a Digital signal Processor (DSP).

10. A method of providing binary compatibility between a base architecture and migrant architecture of a VLIW architecture comprising the steps of:

executing a base execution mode and a migrant execution mode on a base architecture and a migrant architecture, respectively, on a VLIW architecture;

providing a fetch packet retrieved from a memory, the fetch packet having an operating mode in dependence upon the execution mode at the time the request was made to the memory for the fetch packet;

parsing said base architecture mode and migrant architecture mode fetch packets into execute packets and dispatching those base execute packets to the appropriate base architecture decode of the execute hardware on a shared datapath by both the base and migrant architectures;

dispatching execute packet instructions having a migrant execution mode to a migrant architecture decode on a migrant architecture control circuit;

executing execute packet instructions on execution units of execute hardware said execute hardware having a base architecture decode and a migrant architecture decode for decoding said base architecture instructions and said migrant architecture instructions, respectively, in dependence upon the execution mode of the fetch packet of the instructions being decoded, prior to executing;

choosing, in dependence upon the operating mode of said fetch packet, between the output of said migrant architecture decode and the output of said base architecture decode in a multiplexer having one machine word output;

controlling the execution hardware units with said machine word.

11. The method according to Claim 10, and further comprising choosing between the output of said migrant architecture decode and the output of said base architecture decode input and a no operation instruction

5 12. The method according to Claim 10, and further comprising controlling
registers with said machine word.

13. The method according to Claim 10, and further comprising controlling a
global register file with said machine word, which supplies operands to all hardware
10 execution units and accepts results of all hardware execution units.

14. The structure according to Claim 13, and further comprising controlling local
register files that supply operands to either local execution hardware functional units or
neighbor hardware execution functional units subsequent to said controlling said global
15 register file.

15. The method according Claim 14, and further comprising controlling the
various types of execution hardware that evaluate functions on the operands to produce
the results of said hardware execution units subsequent to controlling said local register
20 files.

16. The method according to Claim 10, and further comprising translating
opcodes to the control signals required to execute the specified instructions on the
execution hardware functional units within the base and migrant architecture decode
25 units.

17. The method according to claim 10, wherein said VLIW architecture is a
Digital Signal Processor (DSP).

30 18. The method according to Claim 10 and further comprising the step of issuing
no-operation instruction from said migrant architecture control circuit, to preserve the
semantics of the instructions in the migrant architecture.

Abstract

A subpipelined translation embodiment providing binary compatibility between current and future generations of DSPs is disclosed. When a fetch packet is retrieved from memory, the entire fetch packet is assigned an operating mode (base instruction set or migrant instruction set) according to the execution mode at the time the request was made to the instruction memory for the fetch packet. The fetch packets from the instruction memory are parsed into execute packets and sorted by execution unit (dispatched) in a datapath shared by both execution modes(base and migrant). Because the fetch packet syntax and execution unit encoding is different between the migrant and base architecture in this case, the two execution modes have separate control logic. Instructions from the dispatch datapath are decoded by either base architecture decode logic or the migrant architecture decode logic, depending on the execution mode bound to the parent fetch packet of the instructions being decoded. Code processed by the migrant and base decode pipelines produces machine words that control the register files and the execution hardware functional units. These machine words are selected with a multiplexer. The choice of the eventual machine word from the multiplexer is governed by the operating mode bound to the fetch packet that produced the machine word and sequencing logic for sub-pipelined execution. The selected machine word controls a global register file, which supplies operands to all hardware execution units and accepts results of all hardware execution units.

TI No 29338
Sheet No 1 Of 1PATENT SKETCH FORMAttorney Rebecca H. Capstone Lake

migrant: 1 mul a2,a3,a4
 || and a1,a2,a4
migrant: 2 mul a2,a3,a4
 || or a4,a5,a6
migrant: 3 add a4,a1,a1
 || mul a2,a3,a4
migrant: 4 add a4,a1,a1
 || mul a2,a3,a4
migrant: 5 add a4,a1,a1
 || mul a2,a3,a4

Figure 1

base: 1 mul r2,r3,r4
 || and r1,r2,r4
base: 2 NOP ; inserted NOP
base: 3 NOP ; inserted NOP
base: 4 NOP ; inserted NOP
base: 5 mul r2,r3,r4
 || or r4,r5,r6
base: 6 NOP
base: 7 NOP
base: 8 NOP
base: 9 add r4,r1,r1
 || mul r2,r3,r4
base: 10 NOP
base: 11 NOP
base: 12 NOP
base: 13 add r4,r1,r1
 || mul r2,r3,r4
base: 14 NOP
base: 15 NOP
base: 16 NOP
base: 17 add r4,r1,r1
 || mul r2,r3,r4
base: 18 NOP
base: 19 NOP
base: 20 NOP

Figure 2

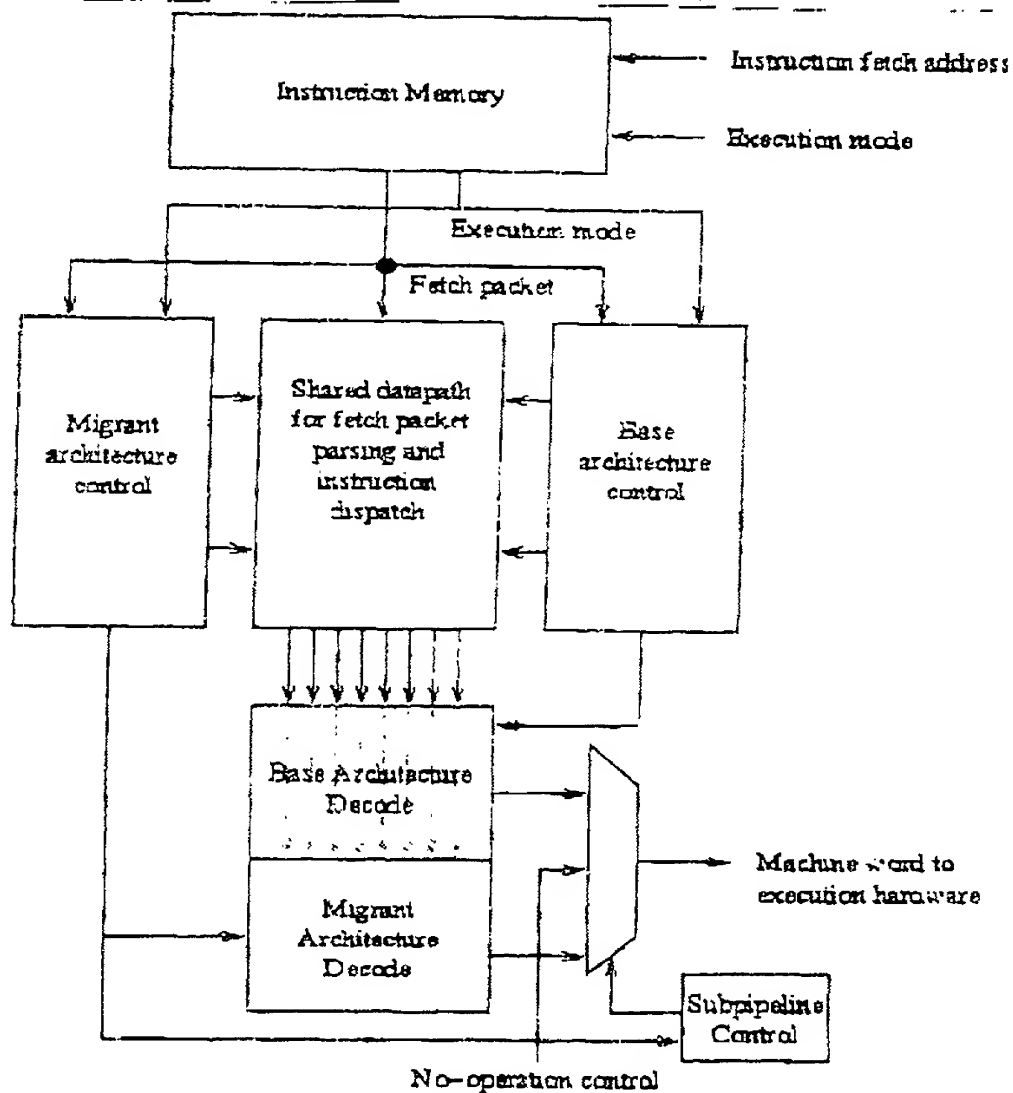
T I No 1111
Sheet No 1 Of 2PATENT SKETCH FORMAttorney Rebecca M. D'Silva, Esq.

Figure 3